ELECTRONICS FOR HIGH-ENERGY PHYSICS EXPERIMENTS IN THE GIGA-SCALE INTEGRATION ERA

K. Shenai and E. McShane, UIC, Chicago, IL USA (Email: shenai@eecs.uic.edu)

Abstract

VLSI technology is being driven to giga-scale levels of integration with IC minimum feature dimensions approaching atomic scales. System-level integration is now pursued as critical in major commercial applications including wireless communication, computing, and multimedia. On-chip signal integrity, noise, and electromagnetic compliance (EMC) are becoming "show-stoppers" in addition to escalating wafer costs. This paper will identify major technology developments and applications in the commercial market and discuss how the high-energy physics community can leverage these advances in years to come. The role of the university research will be discussed as well as new opportunities for collaborative efforts.

1. OVERVIEW AND INTRODUCTION

Functional integration is becoming the dominant industry driver, as systems realize the benefits of reduced IC count, shorter interconnect lengths, and lower power dissipation. System integration, which began with digital functions such as processors and memories, is now being extended to include analog signal processing (ADCs and DACs), RF telecommunications, power generation and conditioning, and pixel sensors. Table 1 lists some of the requirements of these integrated systems from the perspective of HEP experiments.

Table 1: Performance Requirements

System	Performance
Computing bandwidth	> 4 GB/sec
Embedded memory	> 16 MB
Signal processing sensitivity	<< 1 μV
Communications	> 40 MB/sec
Low bit error rate	$< 10^{-9}$
Power conditioning	< 1% ripple
Power levels	> 5 kV
Radiation hardness	> 10 Mrad
Temperature ruggedness	> 150 °C

1.1 Bulk and SOI CMOS

Two flavors of CMOS, bulk and SOI, are available for systems integration, each having specific advantages and limitations. Figure 1 compares the cross sections of a device implemented in bulk and SOI.

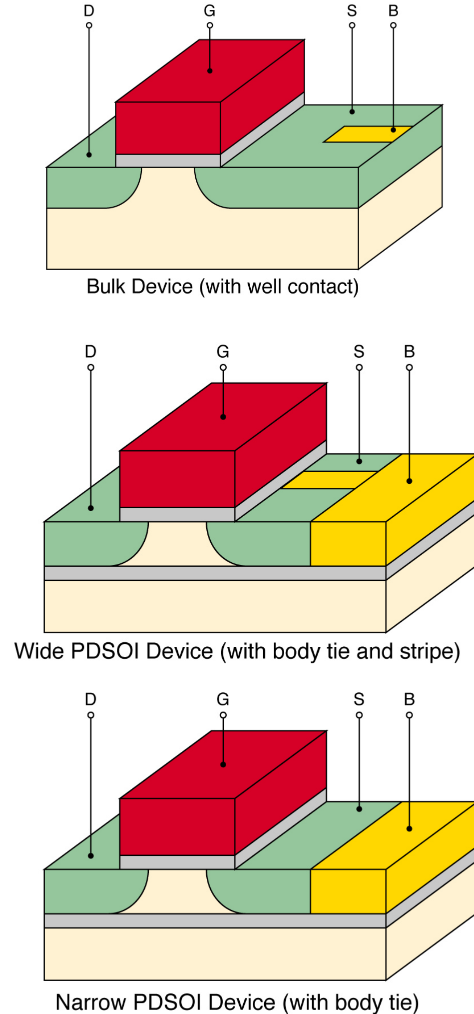

Figure 1: Comparison of bulk and SOI devices.

Bulk has typically better radiation-hardness considering it has no buried oxide with its associated interface traps. SOI, however, offers better isolation given that the buried oxide prevents device crosstalk. Within SOI, both partially and fully depleted implementations are possible. Although fully-depleted SOI provides better electrical performance, the top-gate threshold voltage is linked to the buried oxide potential and is therefore more susceptible to long-term irradiation than PD-SOI.

SOI, despite somewhat poorer radiation tolerance, is an ideal candidate for systems integration. The dielectric isolation permits the co-fabrication of many different

systems, in which the performance of each can be uniquely optimized. Samples are shown in Fig. 2 of SOI implementations of merged power and logic devices, pixel sensing, and combined mixed-signal and RF functions.

Delivering these integrated systems, however, requires the solution to many challenges in signal integrity, noise, EMI, technology optimization, and testability. These issues will be discussed throughout this paper.

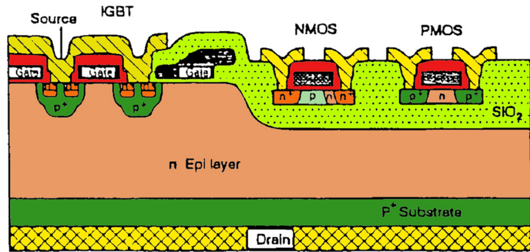

(a)

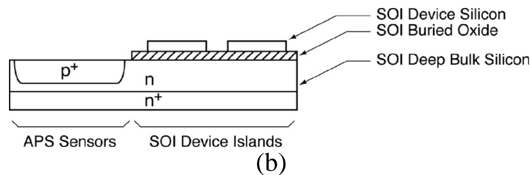

(b)

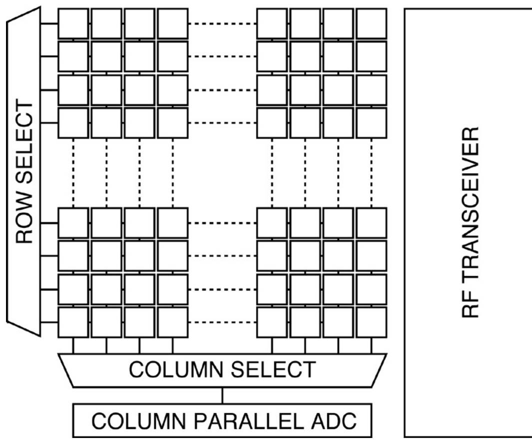

(c)

Figure 2: SOI implementations of (a) merged power and logic devices, (b) pixel sensing, and (c) combined mixed-signal and RF functions.

1.2 Wide Bandgap Semiconductors

In addition to silicon technology, wide bandgap semiconductors such as SiC and GaN are attracting some commercial attention for harsh environments. Although they promise new levels of temperature and radiation robustness, each is immature with numerous material and device challenges remaining. Neither will compete with Si in functional integration applications until these fundamental limitations are overcome. Figure 3 compares the bandgap of three common semiconductors.

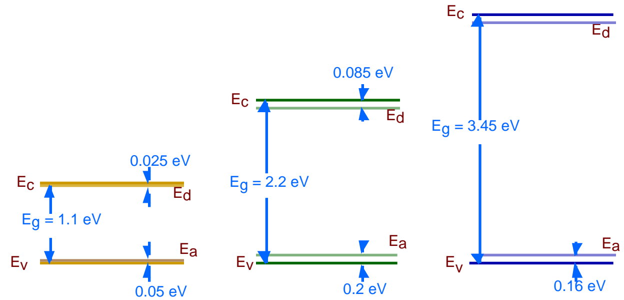

Figure 3: Comparison of wide bandgap semiconductors (Si, SiC, GaN).

1.3 Technology Observations

Despite the wealth of available technologies for commercial and research purposes, semiconductor technologies remain expensive and designs are becoming more sensitive to process variation. Figure 4 identifies some of the trends in IC manufacturing. Complexity has grown commensurately with lithography scaling. Unfortunately, the modeling of these devices has not progressed as effectively. Short-channel effects, gate oxide tunneling, and hot carriers already challenge accepted models. As devices progress below 0.1 μm additional quantum characteristics will demand further characterization and theory.

Apart from purely commercial applications in which power dissipation, form factor, and performance are the key criteria, many of the environmental factors present in HEP experiments are not modeled. For instance, radiation hardness and SEE behavior are not common traits tested by commercial manufacturers. This additional shortcoming is another roadblock to the widespread adoption by the HEP community of off-the-shelf components.

Functionally integrated systems are appearing in consumer markets to meet these demands for compact electronics. But, as device scaling continues to push the envelope toward low-voltage and low-power operation many second-order effects begin to dominate such as EMI, coupling, signal integrity, and non-uniform internal heating. Functionally integrated systems, however, demand additional attention to global signal coupling and isolation, architectural optimization, and power vs. performance tradeoffs.

In following sections the current state-of-the-art in several integrated systems is presented. They show the trend in performance and power metrics with an indication of the COTS suitability for HEP electronics. Many systems in HEP and commercial products have similar performance requirements. The comparison neglects the differences (primarily in reliability and environmental ruggedness) and focuses on performance and efficiency of the electronics.

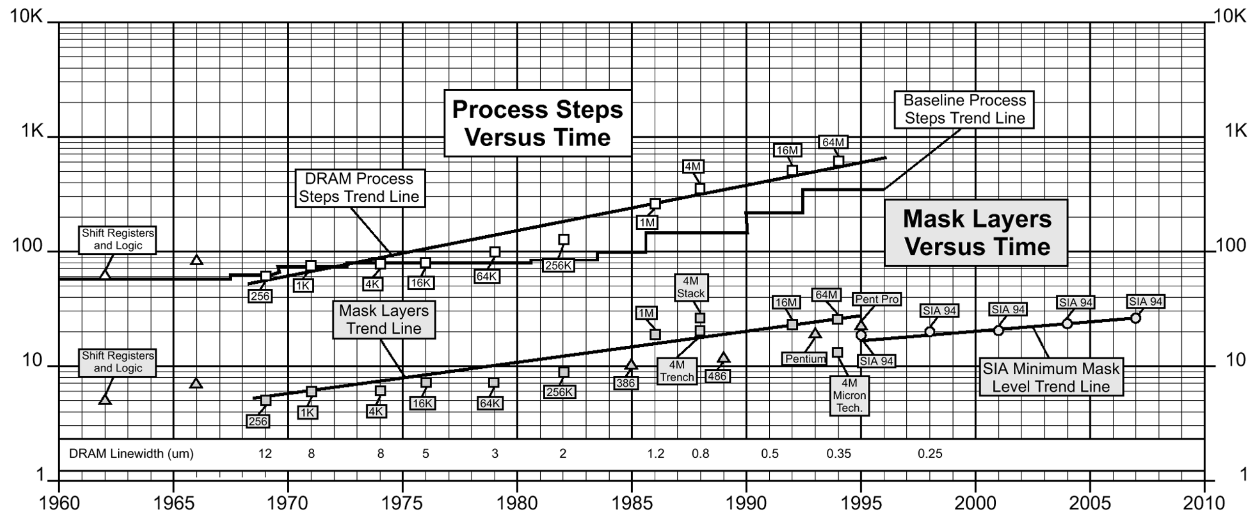

Figure 4: Trends in IC manufacturing complexity with a five-year projection (from Semiconductor Consulting Services, www.semiconsulting.com).

1.4 HEP and Commercial System Similarities

As listed previously in Table 1, many electronics systems comprise the hardware for HEP experiments. Although no commercial application combines the same functional requirements, trends in industry toward functional integration are producing several single-chip solutions to electronics subsystems. A summary of these subsystems is shown in Fig. 5.

Each bubble represents a key market segment. To accommodate new products, particularly in wireless services and multimedia, crossover integration is appearing. For instance RF systems may incorporate digital circuits for baseband processing or a mixed-signal IC may include power management.

The systems selected for review are drawn from this figure: processor/memory, data bus, signal converters, RF, and power management and generation.

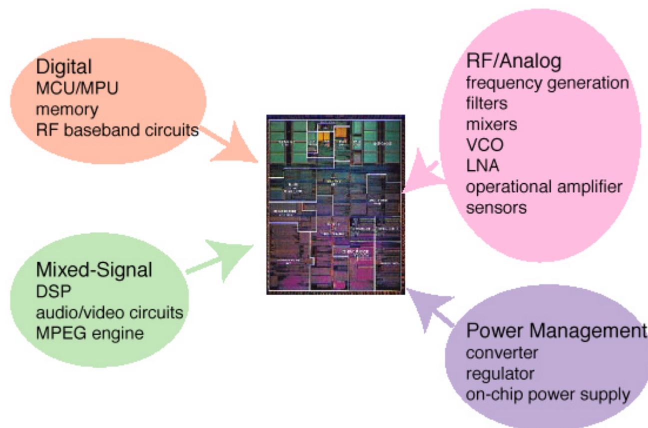

Figure 5: Common systems undergoing functional integration to single-chip products.

2. DIGITAL SYSTEMS

The performance of digital electronics is primarily determined by the processor to memory bandwidth. A typical system block diagram is shown in Fig. 6. It reveals that internal bandwidth far exceeds the capacity of external storage and I/O to supply data. Despite trends to integrate memory closer to the processor, as shown in Fig. 7, bandwidth still limits performance. Two solutions are under investigation for removing the bottleneck. The first involves placing a small set of logic gates into a full memory (DRAM) process: processor-in-memory. The second inserts a small DRAM macrocell into a microprocessor technology: embedded DRAM.

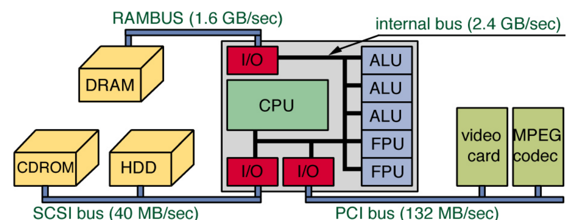

Figure 6: Typical digital system block diagram.

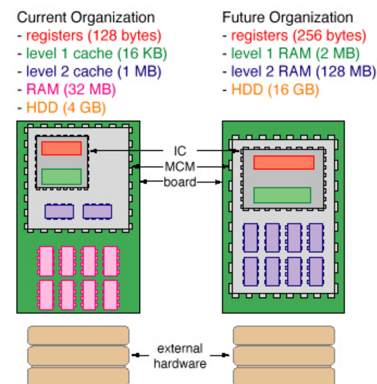

Figure 7: Trends in memory integration.

2.1 Processor-in-Memory

Processor-in-memory (PIM) systems currently integrate on the order of 100K to 500K gates as shown in Fig. 8. This architecture has been most successful for large parallel systems in which simple arithmetic, logical, or conditional operations are to be performed on mass data. Typical examples include image processing, search routines, and associative memories.

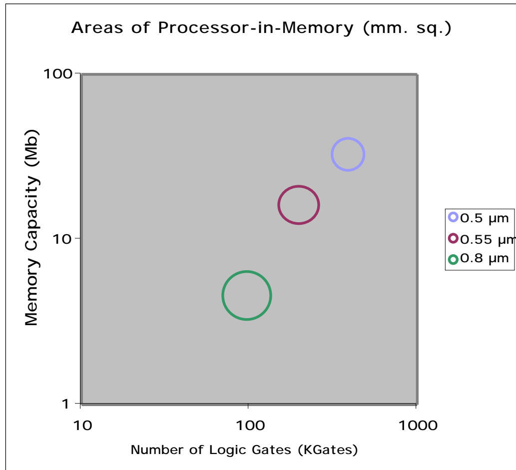

Figure 8: Trends in processor-in-memory (circle area is proportional to die area).

2.2 Embedded Memory

Embedded DRAM macros presently include up to 32Mb by combining smaller 8-Mb macros as shown in Fig. 9. This technology is more process sensitive than PIM since the DRAM is produced in a logic technology. It offers more support for general-purpose computing than PIM and is the likely format for future microprocessors and DSPs. Current applications include complex signal processing that requires high-speed memory intensive algorithms.

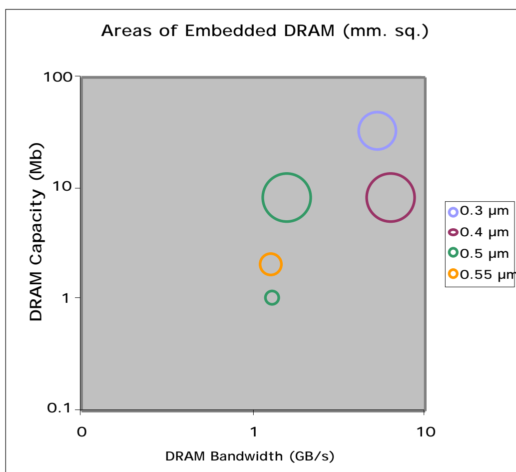

Figure 9: Trends in embedded DRAM (circle area is proportional to die area).

2.3 Memory Partitioning and Optimization

One result of the variety of memory integration approaches and the diversity of logic and memory technologies is a need for new techniques to optimize processor-memory architectures. Shown in Fig. 10 is the resulting performance of five architectures after scaling. We have shown that embedded memory alone does not guarantee best performance [1]. Architectural design remains important, and may scale non-uniformly.

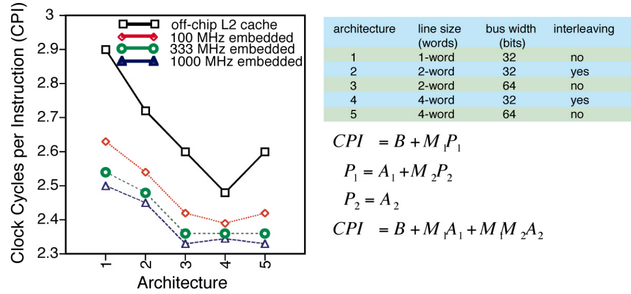

architecture	line size (words)	bus width (bits)	interleaving
1	1-word	32	no
2	2-word	32	yes
3	2-word	64	no
4	4-word	32	yes
5	4-word	64	no

$$CPI = B + M_1 P_1$$

$$P_1 = A_1 + M_2 P_2$$

$$P_2 = A_2$$

$$CPI = B + M_1 A_1 + M_1 M_2 A_2$$

Figure 10: Scaling of memory partitions and interleaving for optimal system performance.

3. ANALOG - DIGITAL CONVERSION

A critical circuit element of many commercial and research electronics systems is the analog-to-digital (ADC) or digital-to-analog (DAC) converter. Although digital electronics dominate high-speed computing and signal processing, the fundamental analog nature of real-time signals requires a translation between the domains. Specifications include power dissipation, sample rate, and bit resolution. These requirements are mutually exclusive, so application-specific designs typically sacrifice one characteristic for another.

Many topologies exist, usually with inherent advantages for either resolution or sample rate. The comparisons below do not distinguish between topology.

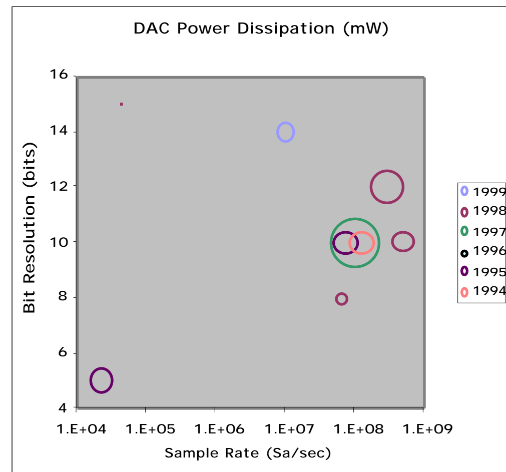

Figure 11: Trends in DACs (circle area is proportional to power dissipation).

3.1 Digital to Analog Converters

Most recent systems favor 10-12 bit resolution to deliver high-quality audio modulation. A few high-resolution (>15 bit) cases exist, but they have considerably lower sampling rates. Typical applications include audio telecommunications and dynamic contrast or threshold control in pixel imagers.

3.2 Analog to Digital Converters

Figure 12 shows a nearly linear trend from low-res, high-rate to high-res, fast-sampling ADCs. The trend reflects the demand for digital signal processing and the need to provide an optimal balance of rate vs. resolution during digitization. Specific topologies yield high-rate (flash) or high-resolution (-) conversion. Multiple architectures such as parallel, pipelined, and folded exist to improve effective conversion rate or for error correction.

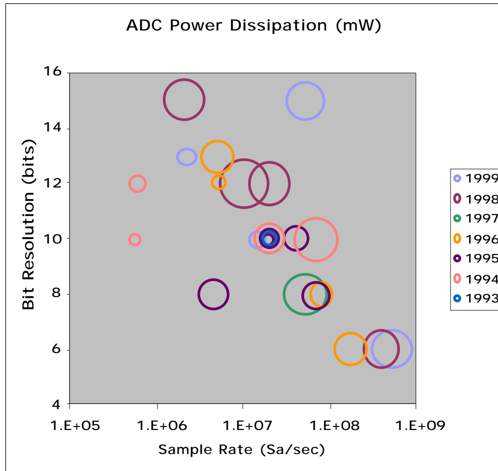

Figure 12: Trends in ADCs (circle area is proportional to power dissipation).

3.3 Wide Bandgap Semiconductors

In high-speed and high-rate converters such as flash ADCs, a significant component of overall power loss is introduced by the auto-zero operation. Another component is the DC current flow through differential amplifier pairs. Figure 13 shows two current waveforms in a flash ADC before and after power minimization.

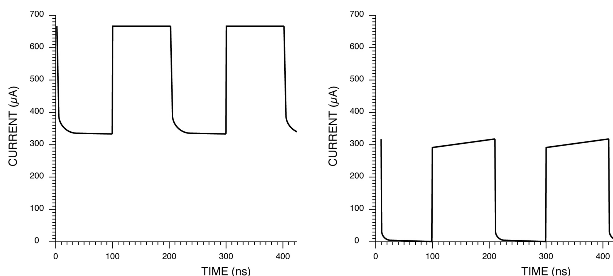

Figure 13: Current waveforms in a flash ADC before and after power-reduction techniques.

We have designed a flash ADC based on dynamic power reduction techniques that achieves a rate of 30-MSa/sec of 500-µV with a per-comparator power loss of 6.25 mW. Alternatively, it can be tuned to deliver 5-MSa/sec of 50-µV with a per-comparator power loss of 5.75 mW. This design was implemented in 0.8-µm SOI CMOS [2-3].

For a 0.35-µm SOI CMOS we expect performance with >100-MSa/sec of <10-µV with a per-comparator power loss of 1 mW

4. RF WIRELESS TRANSCEIVERS

The importance of RF wireless communications to consumer electronics is well established. For the HEP community, wireless telecom promises to eliminate many of the cabling and harness infrastructure that complicates system design and maintenance. Two areas of research are actively pursued. The first is monolithic RF technology to produce compact, very energy efficient radios. The second is architectures with noise suppression and harmonic cancellation to facilitate multi-channel communications with low bit-error rate.

4.1 Transceiver Topologies

Two common topologies are shown in Fig. 14. The shading in the first indicates the chip level implementation of the design. Filters limit the density of the packaging. Technology specific circuits, like Si bipolar mixers and GaAs power amplifiers limit the density of integration.

In the second case, one mixer/filter stage is eliminated by applying frequency translation directly to/from baseband. This zero-IF (direct conversion) transceiver is more amenable to single-chip implementations. Baseband filtering is achieved with digital processing, which can be integrated with the baseband A/D, D/A translation.

Future architectures are looking at extremely high sample rate ADCs and DACs to elevate the baseband processing closer to the RF carrier frequency.

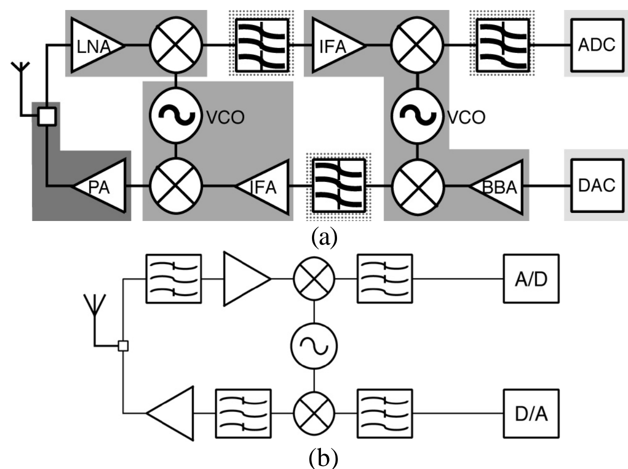

Figure 14: RF transceiver topologies in (a) heterodyne and (b) zero-IF formats.

4.2 Noise Suppression Architectures

Multi-channel communication, especially using new simultaneous access protocols like CDMA, requires broadband linearity with a tight signal envelope. The former is largely a device and packaging matter to ensure correct matching and output characteristics for a full range of power and frequency ratings. The latter requires active cancellation of harmonics (3rd, 5th, and 7th).

Two techniques are shown in Fig. 15 for harmonic cancellation. Both produce out-of-phase components of the high-order harmonics, which when subtracted eliminate the non-linear noise.

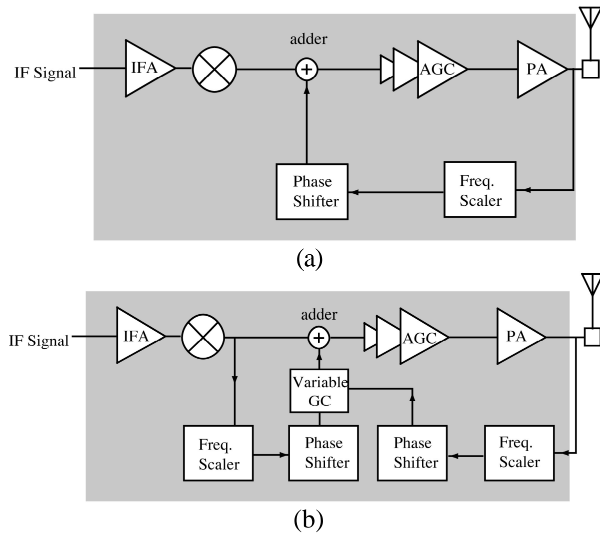

Figure 15: Harmonic cancellation using (a) feedback and (b) bi-directional control.

5. POWER MANAGEMENT

In consumer electronics, power efficiency is often a critical design requirement since many products are battery-operated. On chip power management has several levels of granularity, ranging from global sleep modes to a power-down of specific idle circuits. An enabling technology of power management is on-chip voltage conversion. Many board systems still require 3.3-V or 5-V signaling, but internal IC electronics may prefer <2.5V rails. An on-chip converter can therefore scale the voltage for active operation or suppress the rail voltage for low off-state leakage.

The efficiency of on-chip power electronics is limited by two factors: topology and passives. Unlike high-power electronics, in low-voltage electronics switching and conduction losses through the control devices can quickly lower efficiency below acceptable levels. For example, a diode forward voltage drop of ~1 V represents a significant percentage of the total rail voltage.

Similarly, VLSI passives are generally inferior to discrete elements. Passives in VLSI suffer larger coupling losses and a lack of magnetic materials.

5.1 Converter Topologies

Several integrated converter topologies have been reported, but these typically limit the integration to the digital control electronics; passives are mounted as discrete components. We have developed two generations of fully monolithic DC-DC converter for buck or boost applications [4]. They deliver output voltages as low as 1 V from an input of 3-5 V. Efficiency is over 85% at 100-MHz switching using a VLSI inductor with Q < 5. Its block diagram is shown in Fig. 16.

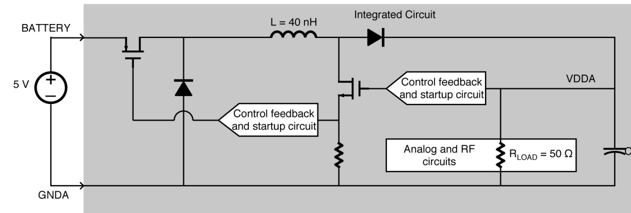

Figure 16: Block diagram of a fully monolithic DC-DC converter.

5.2 Integrated Magnetics

The current manufacturing hurdle for monolithic power electronics is VLSI integrable magnetics. Inductors and transformers are essential elements of power electronics circuits. Their quality factor, Q, directly impacts overall conversion efficiency, as shown in Fig. 17.

The current state-of-the-art in monolithic inductors and transformers is listed in Table 2. Magnetic material strongly improves Q, but its acceptance has been limited because of process incompatibilities with standard CMOS technologies. SOI suspended inductors are an alternative with similarly high Q, although they have been characterized only for RF performance. Their limited current handling may prevent their use in power circuits.

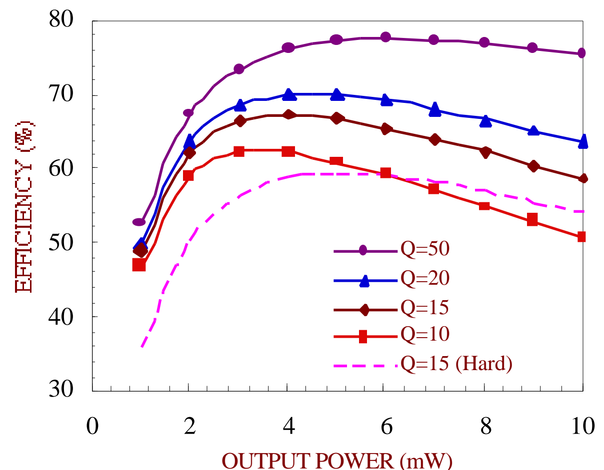

Figure 17: Trends in converter efficiency as a function of inductor quality factor.

Table 2: Monolithic inductor and transformer status.

Quality Factor	Inductance	Comment
VLSI Inductors		
3.5	9.0 nH	spiral
5.7	3.2 nH	hollow core
20.0	1.0 nH	suspended
VLSI transformers		
5.2	9.0 nH	planar interwound
3.7	1.4 nH	planar interwound
Micromagnetics transformers		
12.0	1.0 pH	under 1 MHz

6. POWER GENERATION

Looking beyond on-chip voltage conditioning and level shifting, on-chip power generation promises a new approach for harvesting energy from ambient sources. The simplest system is illustrated by the turbine, which in spinning drives a DC motor. A cross section is shown in Fig. 18. The prime mover fluid can either be a environmental feature (e.g. air or water) or a supplied energy source.

For HEP applications the chief advantages of local power generation are 1) the elimination of magnetic materials, 2) a reduction in lengthy wiring harnesses, 3) the decoupling of noise transients between the load and the energy source, and 4) its compact dimensions. A microturbine implementation is expected to occupy less than 1 cm^2 of surface area.

Figure 19 shows the theoretical performance of a MEMS microturbine. At 7000 rad/sec the generator yields 12 V and 22 W of short circuit power. It can be fabricated with CMOS devices, and SOI offers further options for isolation of MEMS and active devices using the buried dielectric.

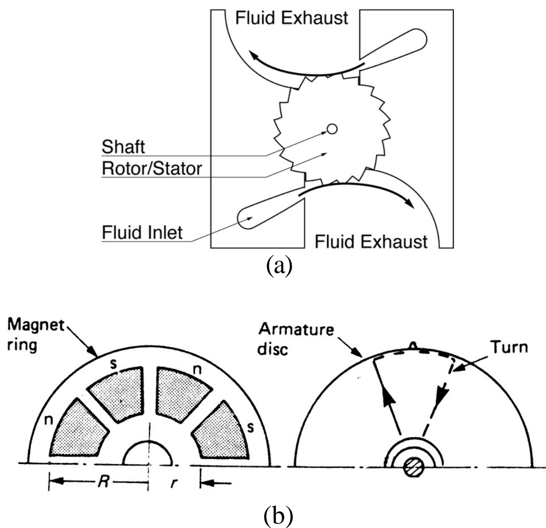

Figure 18: Microturbine showing the top view of its (a) overall structure and (b) rotor and stator patterning.

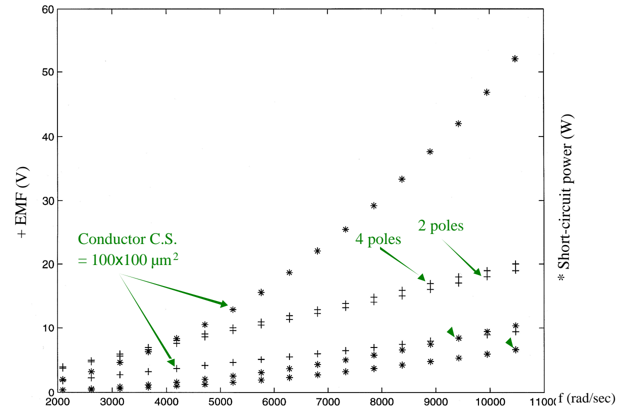

Figure 19: Performance curves of a hypothetical microturbine.

7. RELIABILITY ISSUES

The integration of functional units with widely different dynamic signal characteristics, noise susceptibilities, and activity rates can lead to a significant deterioration of overall signal integrity. Spurious coupling effects also increase the standby power dissipation. As signal rates rise and die sizes shrink, more networks will become sensitive to these transient effects.

Figure 20 illustrates the boundary of transmission line effects for several technologies and physical implementations. As CMOS rise times dip below 10 ps, the transmission-line behavior of digital networks will become common. As Fig. 21 shows, technologies at feature sizes under 0.25-μm with advanced metallization will offer rise times well under 10 ps.

It is therefore necessary to consider the deleterious effects of coupling and impedance mismatch during the entire design phase. This is accomplished through partitioned interconnect routing, the insertion of redundant ground nets, and the insertion of broad ground planes.

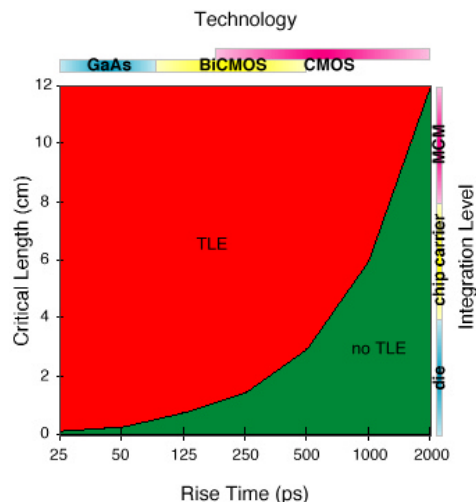

Figure 20: Technology-dependent boundary of transmission line effects.

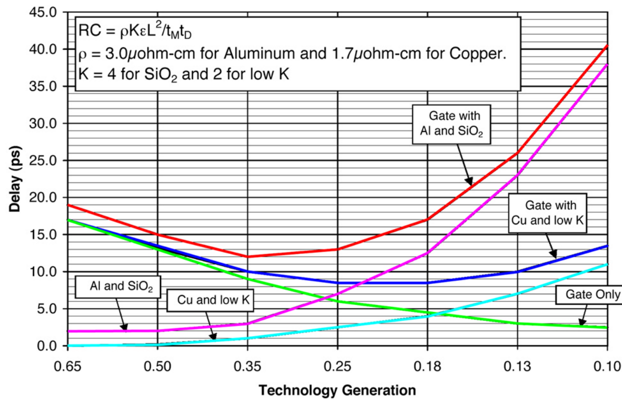

Figure 21: Signal delay as a function of transistor size.

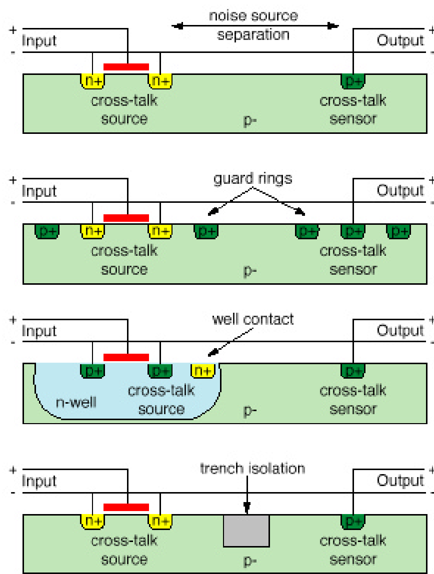

Figure 22: Substrate coupling and suppression.

Figure 22 shows the mechanism of substrate coupling and some techniques for suppressing it. Trench isolation is presently the favored method since it decreases substrate currents without excessively increasing the required area. SOI inherently has no substrate coupling.

8. UNIVERSITY RESEARCH

Correct systems design is becoming more closely linked to IC technology with every new generation. Device models, adequate at large feature sizes, now are challenged to predict comprehensive device behavior under all bias and thermal conditions. Moreover, numerous second-order effects and coupling non-idealities are intruding into nominal system performance. The design of reliable and efficient monolithic systems requires a significant technology expertise. Direct application of COTS technology may fail to address all signal integrity and coupling issues or have insufficient models to predict its performance and ruggedness.

Universities, however, have the resources and collaborations to effectively cultivate the talent for deep submicron functionally integrated system design. Using several research groups at the University of Illinois at Chicago as an example, the argument is presented that universities are ideally suited to act as design and verification centers.

Among the requirements of giga-scale integration are access to a breadth of CAD tools to model systems from the process level to the package/behavioral level. Figure 23 shows the CAD infrastructure resident at UIC. It represents an investment of $50 million and includes software from all major vendors: Analogy, Ansoft, CADENCE, HP EEsof, Silvaco, and Synopsys. It has been integrated into a comprehensive environment for the design and modeling of mixed-signal and RF systems.

Hosting the software is an 8-workstation clustered linked to a multiprocessor server. To support multimillion gate IC design and 10K mesh point finite-element analyses, the network offers over 5 GB of RAM and over 120 GB of permanent storage. The combined computing investment is over $300,000.

Finally, an extensive characterization and instrumentation suite is necessary to extract experimental data. We have a networked mixed-signal/RF testbed. It contains complete digital, analog, and RF experimental instruments and is linked with the computing resource to perform automated measurements and rapid modeling.

These groups at UIC have numerous collaborators within industry and federal research organizations, shown in Fig. 24. As a research entity, they have absorbed the state-of-the-art in design and validation of giga-scale integrated systems. Since industry presently has only a limited interest in radiation-hard electronics for harsh environments, it is necessary to join into teams such as these to address common interests and objectives.

UIC also has access to a variety of commercial and research-oriented foundries, listed in Fig. 25. Deep submicron CMOS offers a contrast to SOI technologies for investigating radiation-tolerance in scaled low-voltage architectures.

9. CONCLUSIONS

Many key components of commercial and HEP-related systems have been investigated with trends outlined toward giga-scale integration. These systems have common performance requirements, except that HEP electronics must survive harsh environments. In addition to anticipated performance limits, deleterious integration effects were discussed. Technology complexity and the lack of industry support for HEP needs limits the application of COTS products. Instead, collaborative research groups led by universities will serve as central design and verification centers by combining multidisciplinary expertise.

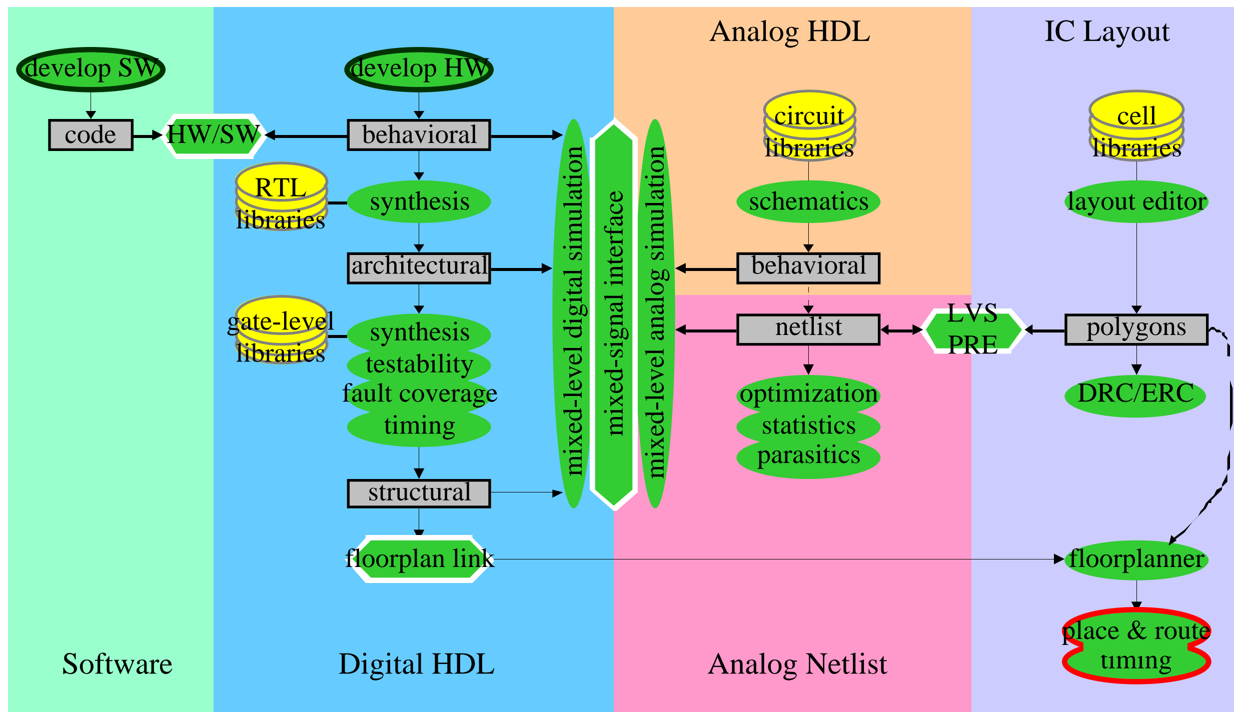

Figure 23: Integrated CAD framework.

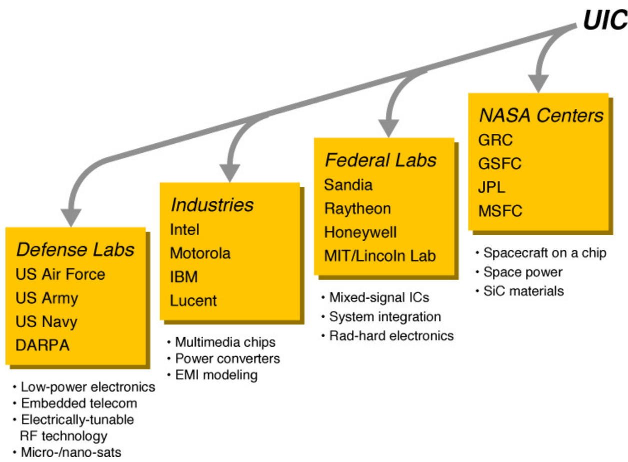

Figure 24: UIC industry and federal collaborators.

Vendor	Feature Size (μm)	Interconnect (Metal/Poly)	Voltage (V)	Analog Extensions	Gate Delay (ps)	Kgates/mm^2
Supertex	2.00	2/2	5	Y	900	1
Peregrine	0.50	3/2	3.3		210	11
AMI	0.50	3/2	5	Y	210	11
	1.20	2/2	5	Y	380	2.8
HP	0.25	5/1	1.8		180	24
	0.35	4/1	3.3		155	20
	0.50	3/1	3.3	Y	210	11
TSMC	0.18	6/1	1.8/3.3		160	28
	0.25	5/1	2.5/3.3	Y	180	24
	0.35	3/2	3.3/5		155	20
	0.35	4/1	3.3/5		155	20

(a)

Vendor	Feature Size (μm)	Interconnect (Metal/Poly)	Voltage (V)	Analog Extensions
Allied Signal	0.50	3/1	5	Y
Honeywell	0.80	3/1	5	Y
	0.35	4/1	3.3	Y

(b)

Figure 25: University access to advanced technologies in (a) bulk and (b) SOI CMOS.

REFERENCES

[1] E. McShane and K. Shenai, "Functionally Integrated Systems on a Chip: Technologies, Architectures, CAD Tools, and Applications," in *IEEE Computer Science Press Innovative Architecture for Future Generation High-Performance Processors and Systems*, pp. 67-75.

[2] E. A. McShane and K. Shenai, "A CMOS Fully-Differential Analog-to-Digital Comparator Cell with Integrated Output Latch and Self-Autozero," Patent Application filed Jan. 1999.

[3] E. A. McShane and K. Shenai, "A Low-Power CMOS Fully-Differential Analog-to-Digital Comparator Cell with Integrated Output Latch and Self-Autozero," Patent Application filed at UIC Jan. 1999.

[4] M. Trivedi and K. Shenai, "Monolithic DC-DC Converter," Patent Application filed at UIC, Jun. 1998.